REMARKS

The present Amendment amends claims 21, 23 and 25, and leaves claims 26 and 27 unchanged. Therefore, the present application has pending claims 21, 23, and 25-27.

35 U.S.C. §103 Rejections

Claims 21-23 and 25-27 stand rejected under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent Application Publication No. 2005/0015460 to Gole et al. ("Gole") in view of U.S. Patent No. 2001/0051955 to Wong. As previously discussed, claim 22 was canceled. Therefore, this rejection regarding claim 22 is rendered moot. This rejection regarding the remaining claims 21, 23 and 25-27 is traversed for the following reasons. Applicants submit that the features of the present invention, as now more clearly recited in claims 21, 23 and 25-27, are not taught or suggested by Gole or Wong, whether taken individually or in combination with each other in the manner suggested by the Examiner. Therefore, Applicants respectfully request the Examiner to reconsider and withdraw this rejection.

Amendments were made to the claims to more clearly describe features of the present invention. Specifically, amendments were made to the claims to more clearly recite that the present invention is directed to a first storage system as recited, for example, in independent claim 21.

The present invention, as recited in claim 21, provides a first storage system connected, via a network, to a computer and a second storage system. The first storage system includes a first storage device which stores data related to a first file system, a first controller which provides the first file system and a second file system to a computer, and a second controller for

controlling input/output operations to/from the second storage system with location of data related to the second file system.

According to the present invention, the second storage system includes a second storage device which stores data related to the second file system and a third controller, connected to the second controller, for controlling the second storage device.

Also according to the present invention, the first controller mounts a root directory of the second file system at a mount point in the first storage system such that the first and second file systems are provided to the computer as a single directory tree.

Furthermore, according to the present invention, the second controller accesses the second storage system with a command representing an area where the data is stored in the second storage device.

Further, according to the present invention, the first storage system is coupled to the second storage system via a storage area network (SAN) and communicates therewith according to a block input/output (I/O) interface, and the first storage system is coupled to the computer via a local area network (LAN) and communicates therewith according to a file I/O interface.

Even further, according to the present invention, each of the first storage device and the second storage device configures a plurality of logical volumes, where when a request to access the second file system is received from the computer, the first controller converts the request into a command for a logical volume of the second storage device, and the second controller sends the command to the second storage system, via the third controller.

Yet even further, according to the present invention, the single directory tree has a total capacity including a capacity of the first storage device and a capacity of the second storage device, and the computer has a transparent single view of the second file system without being aware of whether the second file system resides in the first storage system or the second storage system. The prior art does not teach or suggest all of the above described features.

The above described features of the present invention, as now more clearly recited in the claims, are not taught or suggested by any of the references of record. Specifically, the features are not taught or suggested by either Gole or Wong, whether taken individually or in combination with each other.

Gole teaches a system and method for reliable peer communication in a clustered storage system. However, there is no teaching or suggestion in Gole of the first storage system as recited in claim 21 of the present invention.

Gole discloses a system and method for reliable peer-to-peer communication. The system and method includes a cluster connection manager that acts as a unified resource for opening, closing and maintaining peer-to-peer communication channels between applications executing on storage systems including a storage system cluster. The cluster connection manager communicates with cluster connection manager clients to maintain and manage the communication channels. The cluster connection manager may also be operatively interconnected with a plurality of cluster interconnect devices, thereby enabling failover operation in the event that a cluster interconnect device suffers an error condition.

One feature of the present invention, as recited in claim 21, where the second storage system is of a type different from the first storage system, and when a request to access the second file system is received from the computer, the first controller converts the request into a command for a logical volume of the second storage device, and the second controller sends the command to the second storage system, via the third controller. Gole does not disclose this feature.

For example, the access path of the present invention is quite different from the access path of Gole. In the present invention, as shown in Fig. 13, the computer 400 requests access to the second file system 5300, which is stored in the second storage device 5200 of the second storage system 500. In this way, the request for access is routed from the computer 400, to the first storage system 1, and then to the second storage system 500, for access to the second file system 5300.

This access path taught by the present invention is not the same as the access path taught by Gole. As shown in Fig. 1, and as described in paragraph [0027], Gole discloses a red storage system 200a and a blue storage system 200b, where both the red and blue storage systems have an A loop and a B loop. During normal cluster operation, the storage system that is connected to a disk shelf via the disk shelf's A loop is the "owner" of the disk shelf, and is primarily responsible for servicing data requests directed to blocks on volumes contained on that disk shelf. Thus, the red storage system owns red disk shelf 112 and is primarily responsible for servicing data access requests for data contained on that disk shelf. Similarly, the blue storage system is primarily responsible for the blue disk shelf 114. Accordingly, in

normal cluster operation, the request for access is routed either from the client computer, to the red storage system, and then to the red disk shelf 112, or is routed from the client computer, to the blue storage system, and then to the blue disk shelf 114. During this normal cluster operation, the request for access is not routed from the client computer, to the red storage system, and then to the blue storage system, as in the present invention.

Furthermore, as described in paragraph [0027], when operating as a storage system cluster, each of the red and blue storages systems is configured to take over and assume data handling capabilities for the other disk shelf in the cluster 130 via the disk shelf's B port. Accordingly, with reference to Fig. 1, if the client computer issues a request for access to the blue disk shelf 114, and the A loop of the blue storage system is inoperative, then the request is routed from the client computer to the blue storage system, to the red storage system (via the cluster interconnect 110), and then to the blue disk shelf 114 (via the B loop). In this way, the access path goes from the computer, to the second storage system, and then to the second storage device, which is not the same as the present invention.

Likewise, with further reference to Fig. 1, if the client computer issues a request for access to the red disk shelf 114, and the A loop of the red storage system is inoperative, then the request is routed from the client computer to the red storage system, to the blue storage system (via the cluster interconnect 110), and then to the red disk shelf 114 (via the B loop). In this way, the access path goes from the computer, to the first storage system, and then to the first storage device, which is not the same as the present invention.

By way of further example, Gole does not teach where the first controller converts the request into a command for a logical volume of the second storage device, and the second controller sends the command to the second storage system. As described in paragraph [0031] and [0033] of U.S. Patent Application Publication No. 2005/0182900 of the present application, the computer is connected to the first storage system through a file I/O interface, and the first storage system is connected to the second storage system through a block I/O interface. As further described in paragraph [0111], the present invention provides where when the CHD0 (1110) accesses the heterogeneous storage 500, the CHD0 (1110) is required to use addresses which can be recognized by the disk array controller FCTLx510x of the heterogeneous storage 500. Thus, the CHD0 (1110) accesses the heterogeneous storage 500 using the SLUN (in the volume management table 131 in FIG. 11A) corresponding to the VDEV which stores the data. In this way, the first storage system and the second storage system are accessed using different communications protocols, and accordingly, the first controller of the first storage system converts the request into a command that is recognized by the third controller of the second storage system.

Gole does not teach this feature. To the contrary, in Gole, the red storage system 200a and the blue storage system 200b, are not different types of storage systems, and do not use different communication protocols. Accordingly, it is not necessary to convert the request received from the computer into a command that is recognized by the third controller of the second storage system, as in the present invention, where the second storage system is of a type different from the first storage system (see, e.g.,

paragraph [0039] of U.S. Patent Application Publication No. 2005/0182900 of the present application).

Another feature of the present invention, as recited in claim 21, includes where the single directory tree has a total capacity including a capacity of the first storage device and a capacity of the second storage device, and the computer has a transparent single view of the second file system without being aware of whether the second file system resides in the first storage system or the second storage system. Gole does not disclose this feature.

In the present invention, as described in paragraph [0125] of U.S. Patent Application 2005/0182900 of the present application, a file system PFSx created on a logical device PLDEVx within the storage 1 is combined with a file system SFSx created on a logical device SLDEVx defined in the heterogeneous storage 500 to build a file system PFSx which comprises a single directory tree that has a capacity too large to be built only with the internal disk capacity. With this configuration, the host can use a large scaled file system with a transparent single view without being conscious of whether a file system resides in the storage 1 or in the heterogeneous storage 500. Gole does not disclose this feature.

Therefore, Gole fails to teach or suggest "wherein the second storage system is of a type different from the first storage system, and when a request to access the second file system is received from the computer, the first controller converts the request into a command for a logical volume of the second storage device, and the second controller sends the command to the second storage system, via the third controller" as recited in claim 21.

Furthermore, Gole fails to teach or suggest "wherein the single directory tree has a total capacity including a capacity of the first storage device and a capacity of the second storage device, and the computer has a transparent single view of the second file system without being aware of whether the second file system resides in the first storage system or the second storage system" as recited in claim 21.

The above noted deficiencies of Gole are not supplied by any of the other references of record, namely Wong, whether taken individually or in combination with each other. Therefore, combining the teachings of Gole and Wong in the manner suggested by the Examiner still fails to teach or suggest the features of the present invention as now more clearly recited in the claims.

Wong teaches a mirror file system. However, there is no teaching or suggestion in Wong of the first storage system as recited in claim 21 of the present invention.

Wong discloses a mirror file system (MFS), which is a virtual file system that links two or more file systems together and mirrors between them in real time. When the MFS receives updated data from an application, all file systems linked by the MFS are updated in real time. The file systems linked and mirrored through the mirror file system can be a local file system connected to a physical device, or a network file system exported by a remote system on a network. The real-time mirroring mechanism provided by the MFS is transparent to user applications. The system administrator first sets up the mirroring mechanism by linking a file system to another file system on a single directory through an MFS mounting protocol. These two file systems and their files are linked together and become a mirroring pair. Both copies

are owned by, and under the management of, the MFS. All access to files or directories in both file system go through the MFS. The user applications perform normal file system operation and file/directory operation system calls like open, read, write and close functions from the pathname of either file system. Most of the file operations (such as a read operation) only need to go to one file system under the MFS to get the data. Only when updates occur, such as a write operation, the MFS mechanism ensures that all data updates go to both the file systems. With this mirroring mechanism of the MFS, the files/directories in one file system are mirrored to their mirroring counterparts of another file system in real time. With the MFS technology, a standalone system is able to make multiple copies of data available to the application. In the network environment, multiple servers owning the same data copy can be distributed on the network and mirror the data to each other in real time to provide more efficient and more reliable service to their clients.

One feature of the present invention, as recited in claim 21, includes where the second storage system is of a type different from the first storage system, and when a request to access the second file system is received from the computer, the first controller converts the request into a command for a logical volume of the second storage device, and the second controller sends the command to the second storage system, via the third controller. Wong does not disclose this feature.

To the contrary, Wong is merely directed to a mirror file system, which includes a link between directories. There is no conversion that occurs between file access I/O and block access I/O, as in the present invention.

Another feature of the present invention, as recited in claim 21, includes where the single directory tree has a total capacity including a capacity of the first storage device and a capacity of the second storage device, and the computer has a transparent single view of the second file system without being aware of whether the second file system resides in the first storage system or the second storage system. Wong does not disclose this feature.

Therefore, Wong fails to teach or suggest "wherein the second storage system is of a type different from the first storage system, and when a request to access the second file system is received from the computer, the first controller converts the request into a command for a logical volume of the second storage device, and the second controller sends the command to the second storage system, via the third controller" as recited in claim 21.

Furthermore, Wong fails to teach or suggest "wherein the single directory tree has a total capacity including a capacity of the first storage device and a capacity of the second storage device, and the computer has a transparent single view of the second file system without being aware of whether the second file system resides in the first storage system or the second storage system" as recited in claim 21.

Both Gole and Wong suffer from the same deficiencies, relative to the features of the present invention, as recited in the claims. Therefore, combining the teachings of Gole and Wong in the manner suggested by the Examiner does not render obvious the features of the present invention as now more clearly recited in the claims. Accordingly, reconsideration and

withdrawal of the 35 U.S.C. §103(a) rejection of claims 21, 23 and 25-27 as being unpatentable over Gole in view of Wong are respectfully requested.

The remaining references of record have been studied. Applicants submit that they do not supply any of the deficiencies noted above with respect to the references used in the rejection of claims 21, 23 and 25-27.

In view of the foregoing amendments and remarks, Applicants submit that claims 21, 23 and 25-27 are in condition for allowance. Accordingly, early allowance of claims 21, 23 and 25-27 is respectfully requested.

To the extent necessary, the applicants petition for an extension of time under 37 CFR 1.136. Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, or credit any overpayment of fees, to the deposit account of MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C., Deposit Account No. 50-1417 (500.43772X00).

Respectfully submitted,

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.

CIB/DKM/cp
(703) 684-1120

Carl I. Brundidge
Registration No. 29,621